

ABSTRACT

A system facilitates the distribution and redistribution of chunks of data among multiple servers. The system may identify servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, and failure correlation properties associated with the servers, and place the replicas of the data at the identified servers. The system may also monitor total numbers of replicas of the chunks available in the system, identify chunks that have a total number of replicas below one or more chunk thresholds, assign priorities to the identified chunks, and re-replicate the identified chunks based substantially on the assigned priorities. The system may further monitor utilization of the servers, determine whether to redistribute any of the replicas, select one or more of the replicas to redistribute based on the utilization of the servers, select one or more of the servers to which to move the one or more replicas, and move the one or more replicas to the selected one or more servers.